

REMOTE ACCESS OF THE ILLIAC IV

By K. G. Stevens, Jr.

NASA Ames Research Center

SUMMARY

Until a few years ago, most computational aerodynamic flow simulations were made on general-purpose computers physically located near the researcher wishing the flow simulation. In recent years alternate computer resources have become available. These particular resources are very attractive because they are larger and faster than older computers. This paper will discuss the use of one of these alternate computer resources, namely, the ILLIAC IV.

This discussion will have two major sections. The first section will describe the hardware, that is, the ILLIAC IV, the Illiac system, The Advanced Research Projects Agency (ARPA) computer network, and the IMLAC PDS-1. The second section will trace the execution of the Space Shuttle flow simulation on this hardware. An actual demonstration of this flow simulation will be presented at this conference.

ILLIAC IV

To understand the ILLIAC IV hardware we will look at its four functional parts. Those parts are the control unit, the 64 processing elements, the processing element memories, and the Illiac main memory. (See figure 1 for a diagram of the hardware described below.)

The Control Unit (CU) contains the instruction stack which interprets all instructions, some of which may be completely executed within the CU. Instructions are partially executed and then broadcast to the 64 processing elements; there, the execution is completed by all the processing elements in lock-step. Thus the Illiac operates on up to 64 sets of operands simultaneously. If each operand is viewed as a component of a vector, one may think of the Illiac as a Vector or Array Processor. In addition to managing the Instruction stack, the CU may be thought of as a small self-contained computer. It has four accumulators which are capable of a full set of shifting, bit-setting, and Boolean operations, as well as addition and subtraction. Furthermore, these accumulators may be used as index registers for fetching and storing in the processing elements.

A Processing Element (PE) has six programmable registers, called RGA, RGB, RGS, RGR, RGX, and RGD. The RGA is the accumulator and RGB is its extension; RGS is a scratch register. The remaining registers are somewhat peculiar to the Illiac architecture. RGR is used for inter-PE communications of data. Data may be rotated end-around (data from PE 1 going to PE 64) within the

64 RGRs. RGX acts as an index for intra-PE fetching. This register allows independent fetching depths in each of the PE memories. The RGD contains fault bits and test result bits for that PE. It also contains the bits, called mode bits, which, when set, allow the PE to take part in instructions and, when reset, protect the PE memory as well as RGA, RGS, and RGX from change. The speed of a PE is approximately equal to that of a CDC 6600.

The Processing Element Memories (PEMs) may be thought of in two ways: (1) collectively as 131,072 64-bit words of memory from the CU's point of view, and (2) as a 64 x 2048 matrix of 64-bit words from the point of view of the PEs. In the latter case, each PE is able to access its own column of 2048 words. (Note that the RGX indexing permits the PEs to fetch independently any word within their own column.)

The main memory of the Illiac is logically a 16-million word drum. The drum is divided into 52 bands (tracks) each of which contains 300 Illiac pages (an Illiac page is 1024 64-bit words). The drum may be mapped, that is, data may be stored upon it in predetermined locations, and accessed asynchronously. This enables the programmer to ensure that the data he wishes to fetch are coming under the read/write heads when he needs them. This allows the full billion-bit-per-second transfer rate to be realized during execution. (A detailed description of the Illiac hardware may be found in reference 1.)

THE ILLIAC SYSTEM

The Illiac system includes the ILLIAC IV, the central system, and a B6700. (See figure 2 for a diagram of the hardware described below.) The central system consists of various processors, memories, and devices that interface the Illiac, B6700 and the outside world. The main processor in the central system is currently a PDP-10 running under the TENEX operating system.

File storage is provided by a hierarchy of devices from central memory (PDP-10 memory) to the Unicon laser memory. Files are moved through the storage hierarchy depending on their activity and space availability. The permanent mass storage device is the laser memory which has an on-line capacity of 700 billion bits.

The Burroughs B6700 computer performs utility functions such as assemblies and compilations of GLYPNIR programs. (For a detailed description of this system see reference 2.)

ARPA NETWORK

The ARPA network now has about 50 nodes connected by 50 kilobit lines. (See figure 3 for typical geographical locations.) These nodes fall into one

of two categories, either a Terminal Interface Message Processor (TIP) or an Interface Message Processor (IMP). An IMP can connect up to four computers to the ARPA network. Their basic function is to send to and receive from other IMPs and TIPs strings of bits (either character data or bit data). Figure 4 shows the wide range of computers currently on the network.

A TIP in addition to performing the same functions as an IMP may support dial-up terminals. Almost any terminal is compatible with the TIP. (A complete description of how to use the TIP may be found in reference 3.)

IMLAC PDS-1

IMLAC's PDS-1 consists of a dynamic cathode ray tube (CRT) and a solid-state keyboard controlled by a sixteen-bit 4096-word miniprocessor. This device is capable of emulating a teletype and an IBM 2250 Display Unit.

An acoustic coupler will be used to dial up a TIP. The TIP will then allow the user to access the Illiac system through its IMP. This original connection will be made as a teletype. When graphics data are being transmitted, the PDS-1 will emulate the 2250 Display Unit and display the graphics data on the CRT. When the system stops transmitting graphics data, the PDS-1 will again emulate a teletype so that further instructions may be issued.

THE SHUTTLE CODE

The demonstration problem has been coded by Davy and Reinhardt and resides in the central system memory. This problem consists of computing the inviscid, frozen flow over the first ten meters of a Shuttle Orbiter-like vehicle (based on design version 147). The free-stream conditions correspond to a Mach number of 10 at an altitude of 20 km. The flow field is computed with an angle of attack of 5 degrees to accentuate the relationship between the body shape and the shock shape. A detailed description of the code may be found in reference 4.

THE DEMONSTRATION

The demonstration will make use of most of the hardware and software described above. (See figure 5 for a diagram of the hardware configuration.) The IMLAC PDS-1 will be connected to the MITER TIP via an acoustic coupler and the conventional telephone system. The TIP will connect the IMLAC, via the ARPA network, to the Illiac system's IMP. At this point the IMLAC is logged into the Illiac's PDP-10. To run this version of the shuttle code two parameters are needed. After supplying these two parameters, Illiac execution

may be requested. This request is usually made by submitting a batch job. However, for this demonstration the shuttle code will be run interactively.

The demonstration program will compute two flow fields. The first computation is without canopy, and the fuselage has simply been faired smoothly through the canopy region. The body profile is shown on the screen in solid-line-plot mode. (See figure 6 for a sample CRT picture.) Also displayed on the screen by solid lines are the bow shock locations and the computed body pressure at the leeward symmetry plane as a function of Z, the integration direction.

For the second computation of the flow field, a canopy shape described by a two-parameter (the input parameters) analytical function is added to the fuselage. Results of the flow field as well as the canopy shape are now displayed on the CRT screen in point-plot mode so that they may be contrasted with the previous computation.

While the ILLIAC IV is calculating these two flow fields, the resulting graphics data are stored in Illiac main memory. When the calculations are completed the graphics data are transferred from the Illiac main memory to the central memory. At this point a simple PDP-10 routine is used to transmit the graphics data to the IMLAC where it may be viewed.

This demonstration is intended to show the feasibility of using an advanced computer from a remote location. It also is meant to demonstrate the practicality of using computer flow-field simulations and their graphical representations in solving aerodynamics problems.

REFERENCES

1. Burroughs Corporation:- ILLIAC IV Systems Characteristics and Programming Manual. NASA Contractor Report 2159, 1972.
2. Institute for Advanced Computation: Systems Guide for the ILLIAC IV User. IAC Doc. No. SG-II0000-0000-D, Moffett Field, CA, March 1974.
3. Bolt, Beranek and Newman Inc.: User's Guide to the Terminal IMP. Report No. 2183, December 1974.
4. Davy, W. C.; and Reinhardt, W. A.: Computation of Shuttle Nonequilibrium Flow Fields on a Parallel Processor. Aerodynamic Analyses Requiring Advanced Computers, Part II, NASA SP-347, 1975, pp. 1351-1376.




Figure 1.- ILLIAC IV.

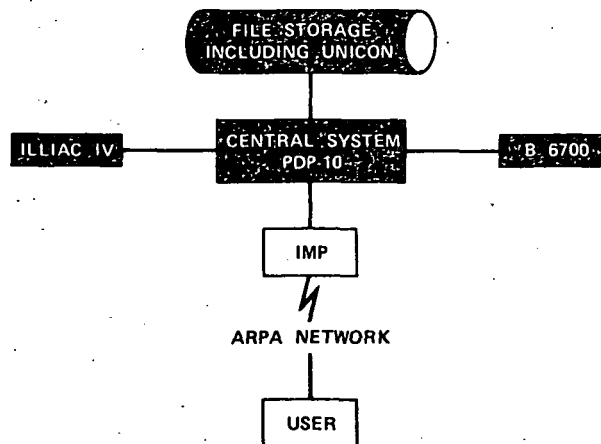


Figure 2.- Illiac system.

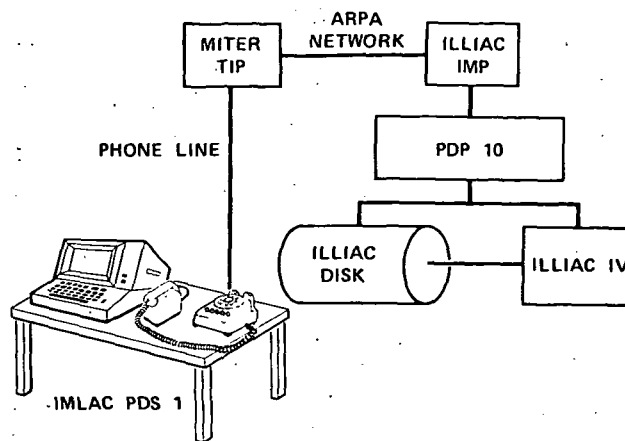


Figure 5.- Demonstration hardware.




Figure 6.- Demonstration frame.